

1 1. A TCP-aware target adapter, for accelerating TCP/IP
2 connections between a plurality of clients and a
3 plurality of servers, the plurality of servers being
4 accessed via an Infiniband fabric, the plurality of
5 clients being accessed via a TCP/IP network, the TCP-
6 aware target adapter comprising:
7 an accelerated connection processor, configured to
8 bridge TCP/IP transactions between the plurality
9 of clients and the plurality of servers, wherein
10 said accelerated connection processor accelerates
11 the TCP/IP connections by prescribing remote
12 direct memory access operations to
13 retrieve/provide transaction data from/to the
14 plurality of servers; and
15 a target channel adapter, coupled to said accelerated
16 connection processor, configured to support
17 Infiniband operations with the plurality of
18 servers, and configured to execute said remote
19 direct memory access operations to
20 retrieve/provide said transaction data.

1 2. The TCP-aware target adapter as recited in claim 1,
 2 wherein said accelerated connection processor
 3 comprises:
 4 a plurality of native network ports, each of said
 5 native network ports communicating with the
 6 plurality of clients in a native network protocol
 7 corresponding to the plurality of clients.

1 3. The TCP-aware target adapter as recited in claim 2,
 2 wherein said native network protocol comprises one of
 3 the following protocols: Ethernet, Wireless Ethernet,
 4 Fiber Distributed Data Interconnect (FDDI), Attached
 5 Resource Computer Network (ARCNET), Synchronous Optical
 6 Network (SONET), Asynchronous Transfer Mode (ATM), and
 7 Token Ring.

1 4. The TCP-aware target adapter as recited in claim 2,
 2 wherein said accelerated connection processor supports
 3 TCP/IP transactions with the plurality of clients by
 4 receiving/transmitting native transactions in
 5 accordance with said native network protocol.

1 5. The TCP-aware target adapter as recited in claim 4,
2 wherein each of a plurality of accelerated TCP/IP
3 connections comprises:
4 a plurality of said remote direct memory access
5 operations between a particular server and said
6 target channel adapter to retrieve/provide
7 particular transaction data from/to said
8 particular server; and
9 corresponding native transactions between said
10 accelerated connection processor and a particular
11 client to provide/retrieve said particular
12 transaction data to/from said particular client.

1 6. The TCP-aware target adapter as recited in claim 5,
2 wherein said accelerated connection processor
3 comprises:
4 a connection correlator, configured to associate TCP/IP
5 connection parameters with a target work queue
6 number for said each of a plurality of accelerated
7 TCP/IP connections.

1 7. The TCP-aware target adapter as recited in claim 6,
2 wherein said TCP/IP connection parameters comprise:
3 source TCP port number, destination TCP port number,
4 source IP address, and destination IP address.

1 8. The TCP-aware target adapter as recited in claim 6,
2 wherein said target work queue number corresponds to a
3 host work queue number within a specific server, said
4 specific server being designated by said accelerated
5 connection processor to support said each of a
6 plurality of accelerated TCP/IP connections with a
7 specific client.

1 9. An apparatus in a server connected to an Infiniband
2 fabric for implementing accelerated TCP/IP connections
3 between the server and clients, the clients being
4 connected to a TCP/IP network, the apparatus
5 comprising:
6 a connection acceleration driver, configured to manage
7 the accelerated TCP/IP connections, wherein said
8 connection acceleration driver designates memory
9 locations within server memory such that
10 transaction data can be retrieved/provided via
11 Infiniband remote direct memory access operations;
12 and
13 a host channel adapter, coupled to said connection
14 acceleration driver, configured to execute
15 Infiniband operations via the Infiniband fabric,
16 and configured to execute direct memory access
17 functions to retrieve/provide said transaction

18 data responsive to said Infiniband remote direct
19 memory access operations.

1 10. The apparatus as recited in claim 9, wherein a
2 particular accelerated TCP/IP connection comprises:
3 a plurality of said remote direct memory access
4 operations between the server and a TCP-aware
5 target adapter to retrieve/provide particular
6 transaction data from/to a particular memory
7 location; and
8 corresponding native transactions between said TCP-
9 aware target adapter and a particular client to
10 provide/retrieve said particular transaction data
11 to/from said particular client.

1 11. The apparatus as recited in claim 9, wherein said
2 connection acceleration driver comprises:
3 native queue logic, configured to interpret a native
4 network protocol corresponding to the clients, and
5 configured to request/receive first Infiniband
6 operations having native TCP/IP transactions
7 to/from the clients that are embedded within
8 Infiniband packets;
9 accelerated queue logic, configured to request second
10 Infiniband operations to establish the accelerated

11 TCP/IP connections, said second Infiniband
12 operations designating said memory locations; and
13 a transport driver interface mux, coupled to said
14 accelerated queue logic, configured to receive
15 said memory locations from application programs,
16 and configured to provide said memory locations to
17 said accelerated queue logic.

1 12. The apparatus as recited in claim 11, wherein said
2 transport driver interface mux is coupled via a
3 transport driver interface to a TCP/IP stack within the
4 server.

1 13. The apparatus as recited in claim 11, wherein said
2 connection acceleration driver further comprises:
3 correlation logic, configured to associate TCP/IP
4 connection parameters with a host work queue
5 number for each of the accelerated TCP/IP
6 connections.

1 14. The apparatus as recited in claim 13, wherein said
2 TCP/IP connection parameters comprise: source TCP port
3 number, destination TCP port number, source IP address,
4 and destination IP address.

1 15. The apparatus as recited in claim 13, wherein said host
2 work queue number corresponds to a target work queue

3 number within a TCP-aware target adapter, said TCP-
4 aware target adapter providing corresponding native
5 transactions to the clients for said each of the
6 accelerated TCP/IP connections.

1 16. An apparatus within a client-server environment for
2 managing an accelerated TCP/IP connection between a
3 server connected to an Infiniband fabric and a client
4 connected to a TCP/IP network, the apparatus
5 comprising:

6 a host driver, for providing a host work queue through
7 which transaction data corresponding to the
8 accelerated TCP/IP connection is
9 transmitted/received via the Infiniband fabric;
10 and

11 a TCP-aware target adapter, coupled to said host
12 driver, for providing a target work queue
13 corresponding to said host work queue, and for
14 executing a remote direct memory access operation
15 to receive/transmit said transaction data via the
16 Infiniband fabric.

1 17. The apparatus as recited in claim 16, wherein said TCP-
2 aware target adapter comprises:
3 a plurality of native network ports, each of said
4 native network ports communicating with TCP/IP

5 clients via a corresponding native network
6 protocol.

1 18. The apparatus as recited in claim 17, wherein said
2 corresponding native network protocol comprises one of
3 the following protocols: Ethernet, Wireless Ethernet,
4 Fiber Distributed Data Interconnect (FDDI), Attached
5 Resource Computer Network (ARCNET), Synchronous Optical
6 Network (SONET), Asynchronous Transfer Mode (ATM), and
7 Token Ring.

1 19. The apparatus as recited in claim 17, wherein said TCP-
2 aware target adapter further comprises:
3 an accelerated connection processor, for supporting
4 TCP/IP transactions with the clients by
5 receiving/transmitting native transactions in
6 accordance with said native network protocol.

1 20. The apparatus as recited in claim 19, wherein said TCP-
2 aware target adapter further comprises:
3 a connection correlator, for associating TCP/IP
4 connection parameters for the accelerated
5 connection with said target work queue.

1 21. The apparatus as recited in claim 20, wherein said host
2 driver comprises:
3 connection correlation logic, for associating said
4 TCP/IP connection parameters for the accelerated
5 connection with said host work queue.

1 22. The apparatus as recited in claim 21, wherein said
2 TCP/IP connection parameters comprise: source TCP port
3 number, destination TCP port number, source IP address,
4 and destination IP address.

1 23. A method for accelerating TCP/IP connections in a
2 client-server environment having clients that are
3 connected to a TCP/IP network and servers that are
4 connected to an Infiniband fabric, the method
5 comprising:
6 a) mapping TCP/IP connection parameters for accelerated
7 connections to corresponding host and target work
8 queue pairs; and
9 b) executing Infiniband remote direct memory access
10 operations to retrieve/transmit data associated
11 with the accelerated connections from/to memory
12 within the servers.

1 24. The method as recited in claim 23, wherein said mapping
2 comprises:

- 3 i) intercepting the TCP/IP connection parameters from
4 requests to send/receive data from/to the servers;
5 and
6 ii) establishing Infiniband connections between the
7 servers and a TCP-aware target adapter.

1 25. The method as recited in claim 24, wherein said
2 executing comprises:

- 3 i) providing the TCP-aware target adapter with memory
4 locations within the servers for
5 transmission/reception of the data;
6 ii) from the TCP-aware target adapter, transmitting the
7 remote direct memory access operations to the
8 servers; and
9 iii) from the servers, providing remote direct memory
10 access responses.

1 26. The method as recited in claim 23, further comprising:
2 c) generating TCP/IP transactions in a native network
3 protocol to provide the data to the clients.

1 27. A method for offloading server TCP/IP processing in a
 2 client-server environment, comprising:
 3 a) utilizing remote direct memory access operations via
 4 an Infiniband fabric to directly access data
 5 from/to server memory, wherein the data is
 6 provided to/from a TCP-aware target adapter, the
 7 TCP-aware target adapter providing native network
 8 ports that connect to clients; and
 9 b) via the TCP-aware target adapter, generating native
 10 network transactions to transfer the data to/from
 11 clients.

1 28. The method as recited in claim 27, wherein said
 2 utilizing comprises.
 3 i) associating TCP/IP connection parameters for a
 4 particular TCP/IP connection with a work queue
 5 within the TCP-aware target adapter; and
 6 ii) issuing remote direct memory access requests to the
 7 work queue.

1 29. The method as recited in claim 28, wherein said
 2 generating comprises:
 3 i) formulating TCP headers, IP headers, and native
 4 network headers for messages to/from the clients
 5 based upon the TCP/IP connection parameters
 6 provided by said associating.

1 30. A TCP-aware target adapter, for accelerating TCP/IP
2 connections between a plurality of clients and a
3 plurality of servers, the plurality of servers being
4 accessed via an Infiniband fabric, the plurality of
5 clients being accessed via a TCP/IP network, the TCP-
6 aware target adapter comprising:
7 an accelerated connection processor, configured to
8 bridge TCP/IP transactions between the plurality
9 of clients and the plurality of servers, wherein
10 said accelerated connection processor accelerates
11 the TCP/IP connections by prescribing remote
12 direct memory access operations to
13 retrieve/provide transaction data from/to the
14 plurality of servers; and
15 a target channel adapter, coupled to said accelerated
16 connection processor, configured to support
17 Infiniband operations with the plurality of
18 servers, and configured to execute said remote
19 direct memory access operations to
20 retrieve/provide said transaction data, and
21 configured to route said transaction data to/from
22 the plurality of clients as embedded payloads
23 within Infiniband packets.

1 31. The TCP-aware target adapter as recited in claim 30,
2 wherein said accelerated connection processor supports
3 TCP/IP transactions with the plurality of clients by
4 formatting and processing native transactions in
5 accordance with a native network protocol corresponding
6 to the plurality of clients.

1 32. The TCP-aware target adapter as recited in claim 31,
2 wherein said accelerated connection processor
3 encapsulates outgoing TCP/IP transactions within
4 Infiniband raw packets for transmission to the
5 plurality of clients.

1 33. The TCP-aware target adapter as recited in claim 32,
2 wherein each of a plurality of accelerated TCP/IP
3 connections comprises:
4 a plurality of said remote direct memory access
5 operations between a particular server and said
6 target channel adapter to retrieve/provide
7 particular transaction data from/to said
8 particular server; and
9 corresponding native transactions between said
10 accelerated connection processor and a particular
11 client to provide/retrieve said particular
12 transaction data to/from said particular client,

13 wherein said corresponding native transactions are
14 encapsulated within Infiniband raw packets.

1 34. The TCP-aware target adapter as recited in claim 33,
2 wherein said accelerated connection processor
3 comprises:
4 a connection correlator, configured to associate TCP/IP
5 connection parameters with a target work queue
6 number for said each of a plurality of accelerated
7 TCP/IP connections.

1 35. The TCP-aware target adapter as recited in claim 34,
2 wherein said TCP/IP connection parameters comprise:
3 source TCP port number, destination TCP port number,
4 source IP address, and destination IP address.

1 36. The TCP-aware target adapter as recited in claim 35,
2 wherein said target work queue number corresponds to a
3 host work queue number within a specific server, said
4 specific server being designated by said accelerated
5 connection processor to support said each of a
6 plurality of accelerated TCP/IP connections with a
7 specific client.

1 37. The TCP-aware target adapter as recited in claim 33,
2 wherein said connection correlator associates native
3 connection parameters with a target work queue number

4 for said each of a plurality of unaccelerated TCP/IP
5 connections.

1 38. The TCP-aware target adapter as recited in claim 37,
2 wherein said native connection parameters comprise:
3 source MAC address and destination MAC address.

1 39. The TCP-aware target adapter as recited in claim 38,
2 wherein said target work queue number corresponds to a
3 host work queue number within a specific server, said
4 specific server being designated by said accelerated
5 connection processor to support said each of a
6 plurality of unaccelerated TCP/IP connections with a
7 specific client.

1 40. An Infiniband-to-native protocol translation apparatus,
2 for routing TCP/IP transactions between a plurality of
3 clients and a plurality of Infiniband devices, the
4 plurality of Infiniband devices being accessed via an
5 Infiniband fabric, the plurality of clients being
6 accessed via a TCP/IP network, the Infiniband-to-native
7 protocol translation apparatus comprising:
8 an unaccelerated connection processor, configured to
9 bridge the TCP/IP transactions between the
10 plurality of clients and the plurality of
11 Infiniband devices by encapsulating/stripping the
12 TCP transactions within/from Infiniband raw

13 packets, said unaccelerated connection processor
 14 comprising:
 15 an unaccelerated connection correlator, for
 16 mapping native addresses to/from Infiniband
 17 local indentifiers and work queue numbers;
 18 and
 19 a target channel adapter, coupled to said unaccelerated
 20 connection processor, configured to
 21 receive/transmit said Infiniband raw packets
 22 from/to the plurality of Infiniband devices.

1 41. The Infiniband-to-native protocol translation apparatus
 2 as recited in claim 40, wherein said native addresses
 3 comprise MAC addresses.

1 42. The Infiniband-to-native protocol translation apparatus
 2 as recited in claim 40, wherein said native addresses
 3 comprise IP addresses.

1 43. The Infiniband-to-native protocol translation apparatus
 2 as recited in claim 40, wherein said Infiniband local
 3 identifiers comprise source local identifier,
 4 destination local identifier, and work queue number.

1 44. The Infiniband-to-native protocol translation apparatus
 2 as recited in claim 43, wherein said Infiniband local
 3 identifiers map said TCP/IP transactions between a

4 particular client and a server connected to an
5 Infiniband fabric.

1 45. The Infiniband-to-native protocol translation apparatus
2 as recited in claim 43, wherein said Infiniband local
3 identifiers map said TCP/IP transactions between a
4 particular client and a TCP-aware target adapter
5 connected to an Infiniband fabric.

109